

# DOEs Advanced Network Initiative (ANI) Testbed Project

Brian Tierney

ESnet

April 12, 2011



# ANI Testbed Overview



## Consists of 3 Phases

- Phase 1: “Tabletop testbed” at LBL (June 2010 to March 2011)
- Phase 2: Move to Long Island MAN when dark fiber is available (April 2011 to September 2011)
- Phase 3: Extend to WAN when 100Gbps available (late 2011)

Funded as a 3 year project

## Capabilities

- Ability to support end-to-end networking, middleware and application experiments, including interoperability testing of multi-vendor 100Gbps network components
- Researchers get “root” access to all devices
- Use Virtual Machine technology to support custom environments
- Detailed monitoring capabilities

# Sample Projects



Examples of the types of projects the current testbed will support include the following:

- Path computation algorithms that incorporate information about hybrid layer 1, 2 and 3 paths, and support 'cut-through' routing
- New transport protocols for high speed networks
- Protection and recovery algorithms
- Automatic classification of large bulk data flows
- New routing protocols
- New network management techniques
- Novel packet processing algorithms
- High-throughput middleware and applications research

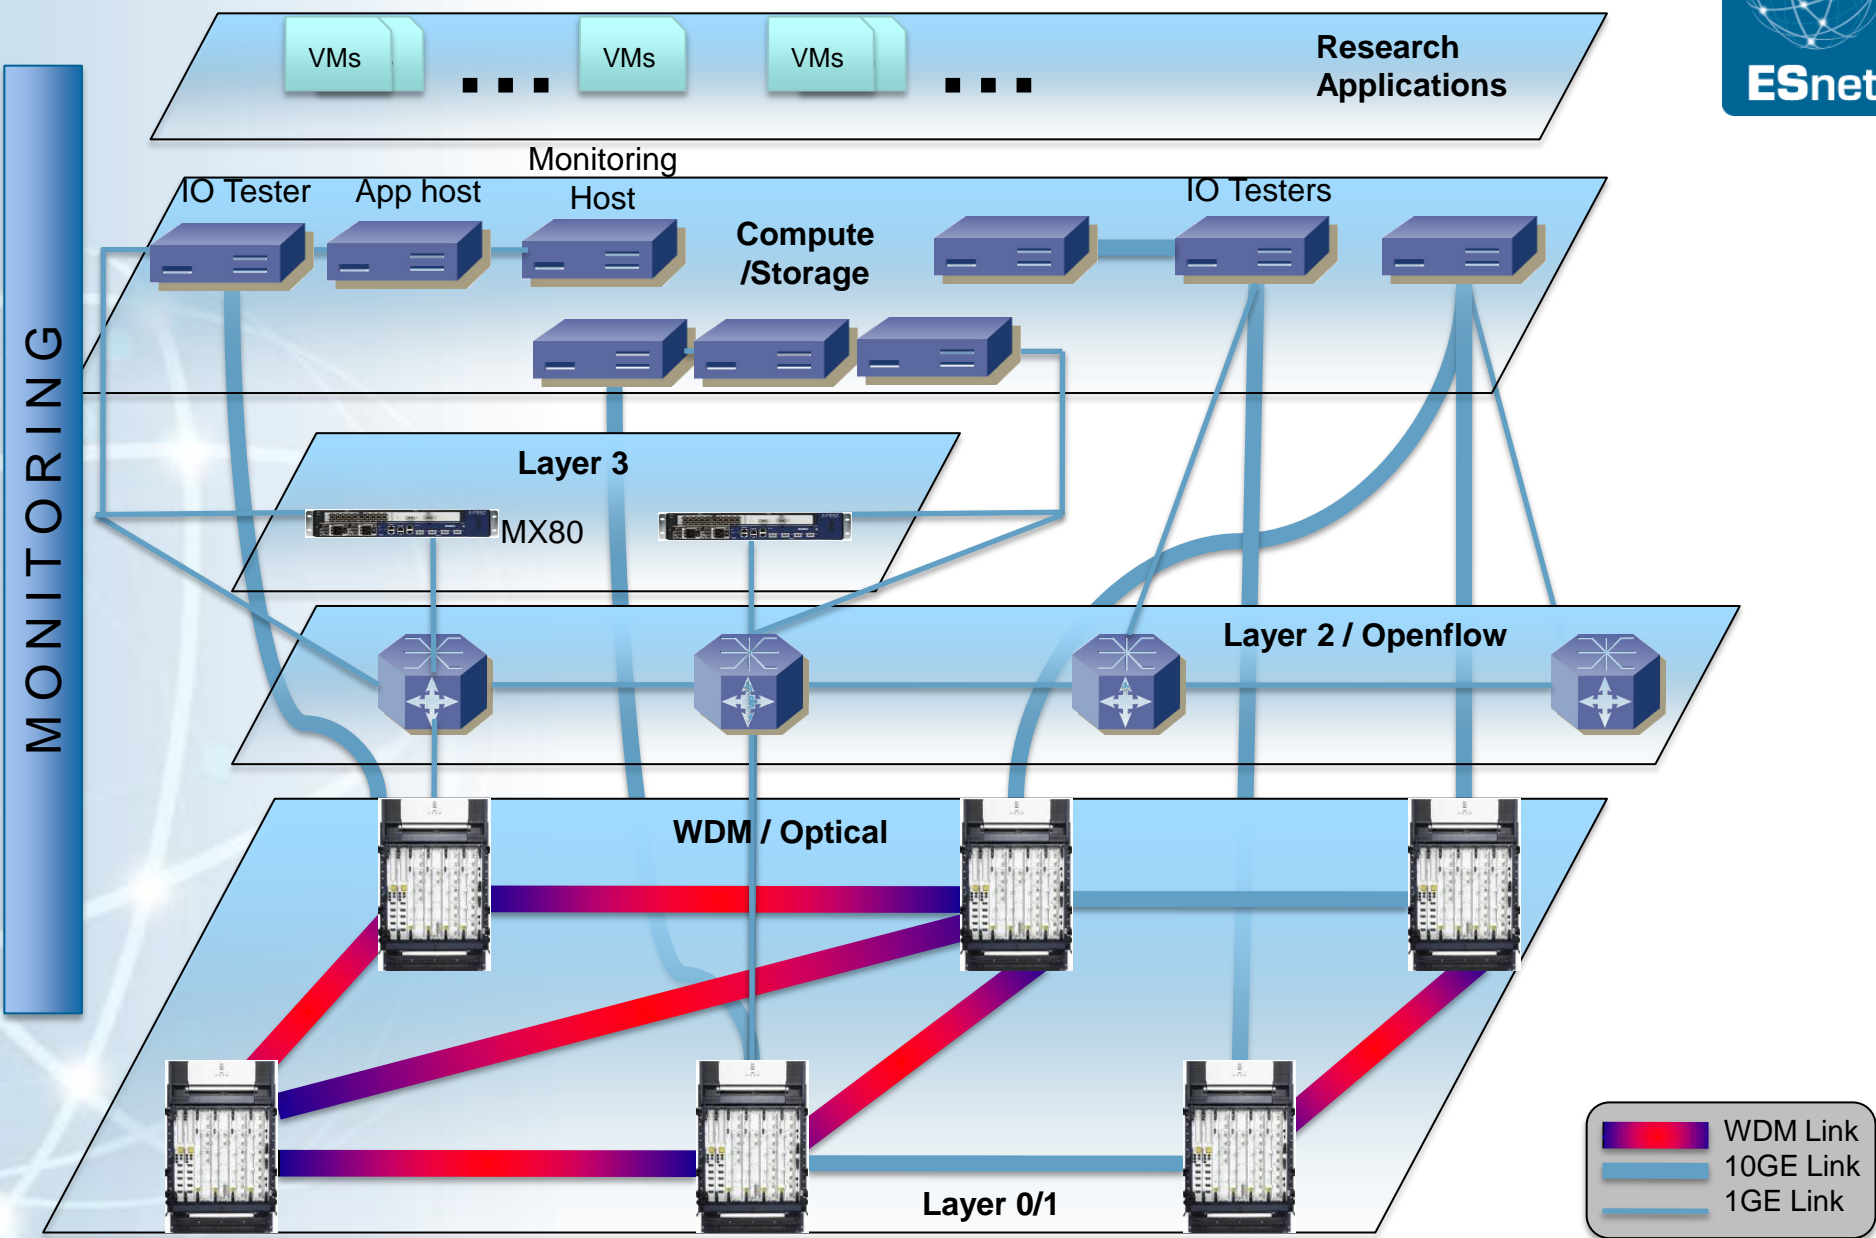
# Network Testbed Components

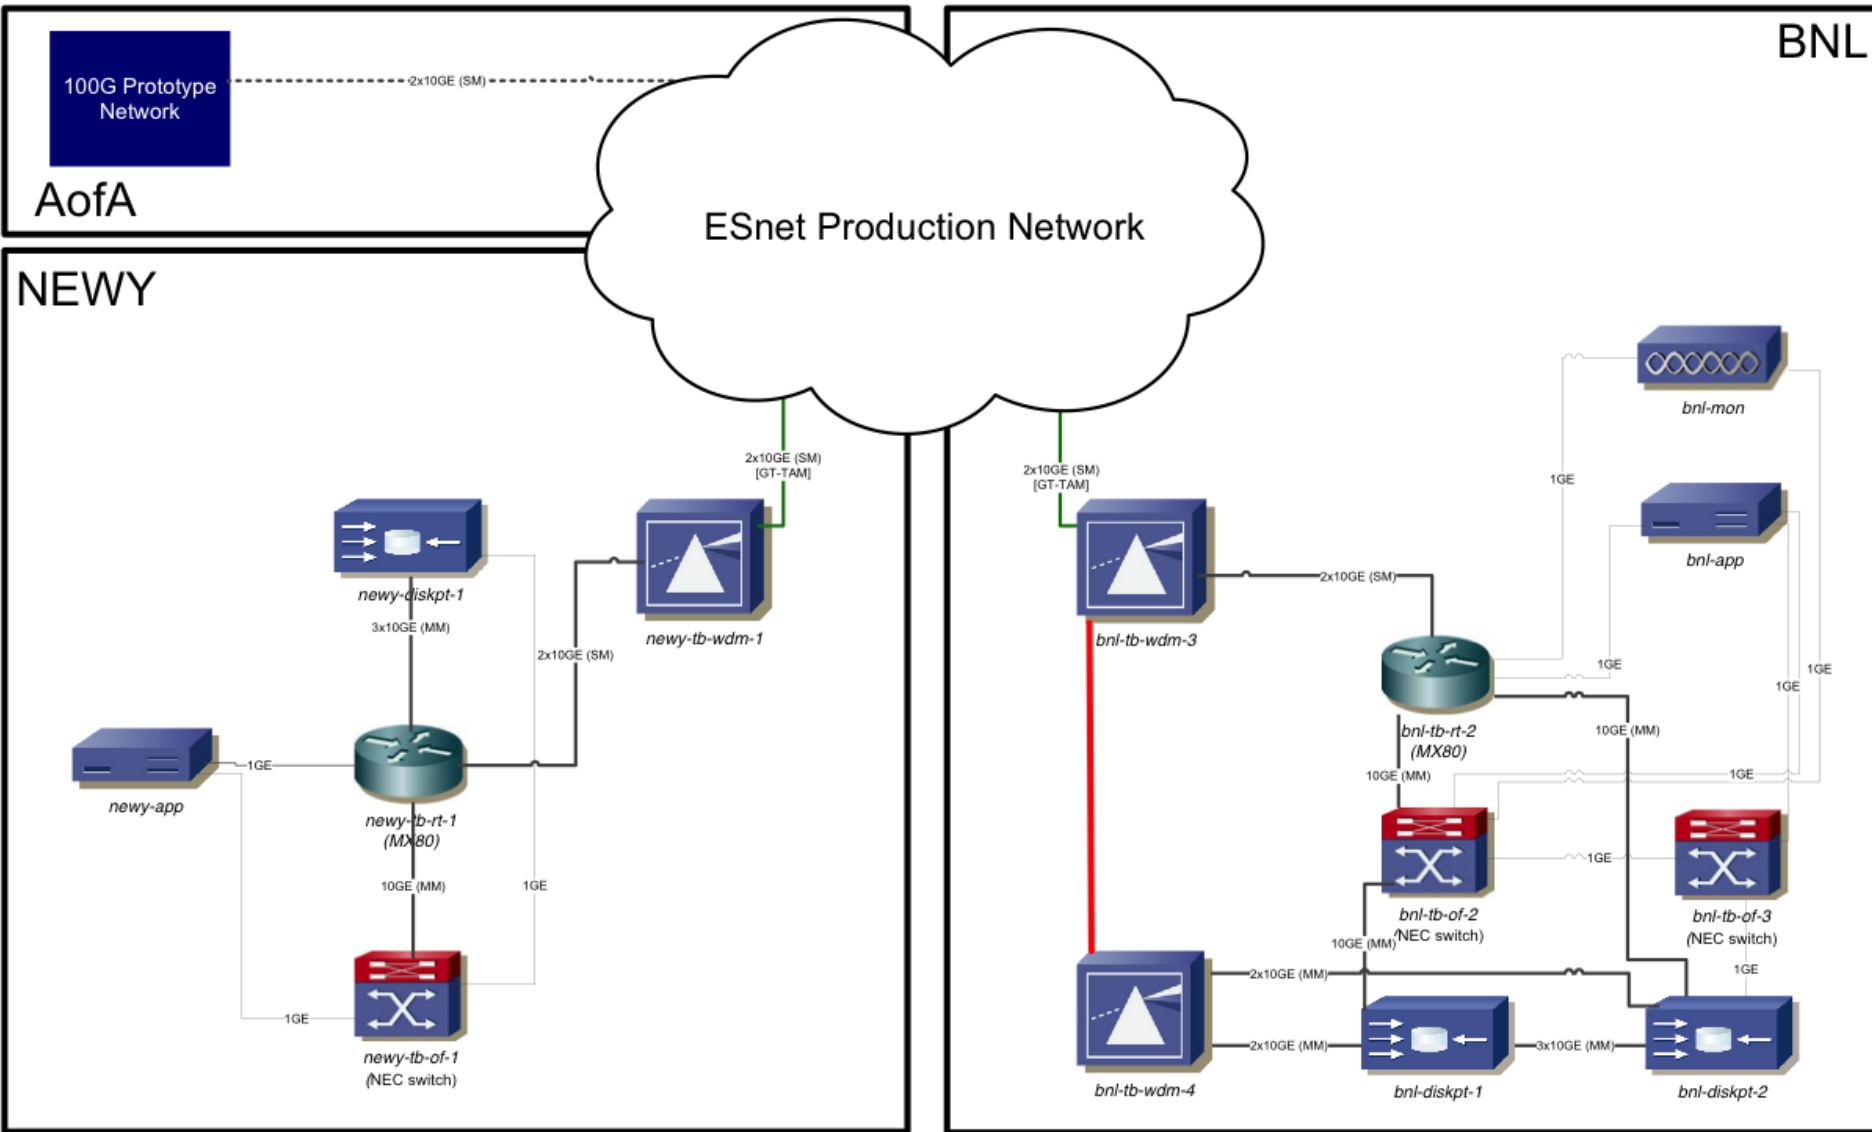


Table Network Testbed consists of:

- DWDM devices (Layer 0-1)
- Layer 2 switches supporting Openflow
- Layer 3 Routers
- Test and measurement hosts
  - Virtual Machine based test environment
  - 4 or 6 x 10G test hosts initially
    - Eventually 40G and 100G from Acadia 100G NIC project
- This configuration will evolve over time

# ANI Testbed: A Layered View





Updated March 22, 2011

**Notes:**

- “App Host”: can be used for researcher application, control plane control software, etc. Can support up to 8 simultaneous VMs
- “I/O Testers” are capable of 15 G disk-to-disk or 35G memory-to-memory
- Other infrastructure not shown: VPN Server, file server (NFS, webdav, svn, etc.)

# Testbed Access

Proposal process to gain access described at:

<https://sites.google.com/a/lbl.gov/ani-testbed/>

Testbed is available to anyone:

- DOE researchers
- Other government agencies
- Industry

Must submit a short proposal to the testbed review committee

- Committee is made up of members from the R&E community and industry

First round of proposals were due Oct 1

- Accepted proposals announced Dec 10, 2010

# Testbed Access (cont.)



Plan is to accept roughly five proposals every 6 month review cycle

- Last round of proposals was due April 1, 2011
- Next round of proposals due Oct 1, 2011

Proposals are reviewed by a review committee with DOE, NSF, University, Industry, and International R&E community members



# Acceptance Criteria



The criteria for selecting the proposals will be based on:

- Quality of proposed research which includes:
  - Clear, focused research topic
  - Creative and original concept
  - Test plan
- Qualifications of the team
- Potential impact of the research on field of networking and DOE SC mission
- Readiness: is the project ready to run experiments right away?
- Value of ANI testbed resources to the research
- Level of support required by ESnet staff

# Current ANI Testbed Research

12 projects have been given access to the testbed so far:

- 6 via direct DOE/ASCR funding

- 6 via testbed proposal process

Types of projects:

- 4: high-speed middleware

- 3: network control plane

- 1: 100Gbps end host hardware

- 1: network flow classification

- 1: TCP congestion control

- 1: security

- 1: energy efficiency

# Sample ANI Testbed Project: Hybrid Network Traffic Engineering Software (HNTES)



PI: Malathi Veeraraghavan, University of Virginia

- Investigating the role of hybrid networking at 100Gbps

Project goal:

- To learn how to optimize a hybrid network comprised of an IP datagram network and a high-speed optical dynamic circuit network
  - Because large-sized flows adversely effect small-sized flows, use machine learning techniques to identify large-sized flows based on size and duration
  - Upon detecting such a flow, HNTES reconfigures the router to redirect packets from this flow to a circuit on a different path

Current status: completed phase I of the software implementation, and demonstrated it at the Oct. 2010 DOE annual review meeting

# ANI Testbed: Phase 3 (October 2011 till ...)



Break into 2 separate testbeds:

- control plane testbed in Long Island
- 100 Gbps application/middleware testbed: NERSC to ANL

Advantages:

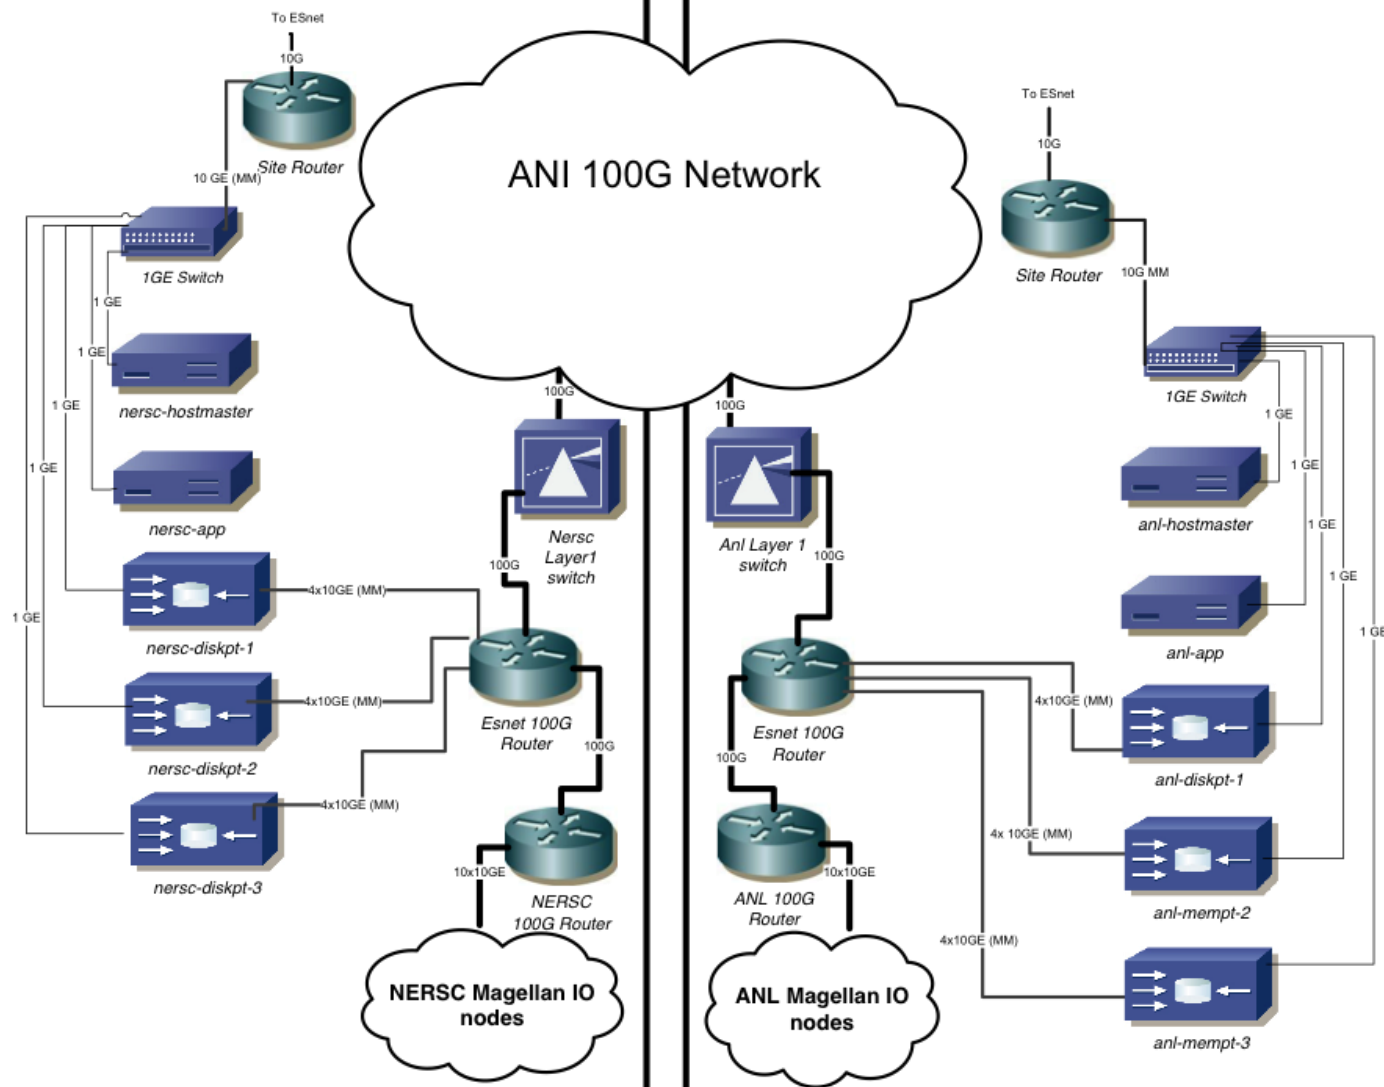
- Leverages the 100G WAVE from NERSC to ALCF
- Easy to modify the configuration, add research equipment, etc
  - NERSC is close to Berkeley, and ESnet staff located at ANL
- Separate control plane and application testing wont interfere with each other

Testbeds will be connected together at 20 Gbps if both parts are needed

# ANI Middleware Testbed Architecture Plan (10-1-2011)

NERSC

ANL

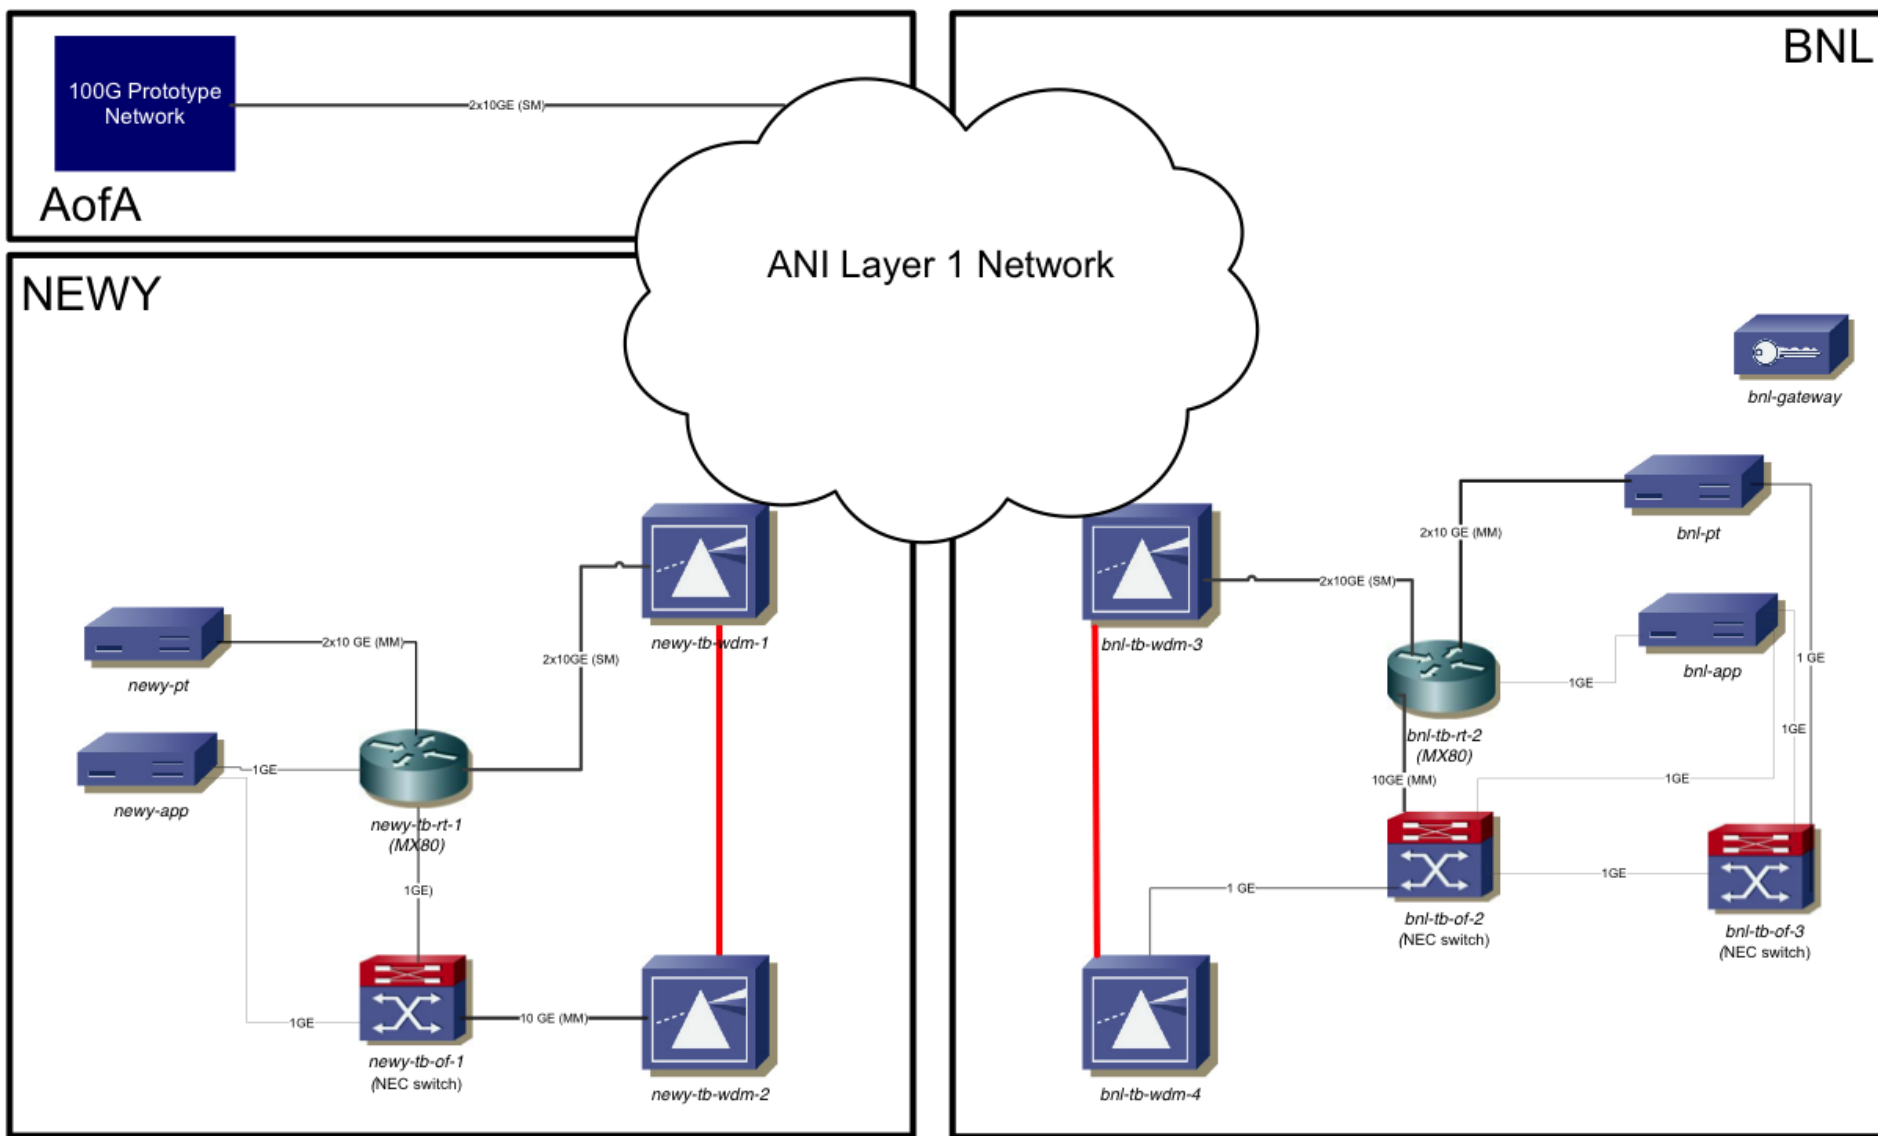


Updated April 3, 2011

## Notes:

- 12x10GE connections at both sides
- Magellan Project resources will be connected 10/2011 to 12/2011

# LIMAN Control Plane Testbed Architecture Plan (starting 10-1-11)



Updated March 22, 2011

## Notes:

- Includes Openflow Testbed
- Includes 2x10Gbps to support combined control plane / application research
- Connects to ANI prototype network at 2x10GE.

# More Information




<http://sites.google.com/a/lbl.gov/ani-testbed/>

<http://100gbs.lbl.gov/>

email: **ani-testbed-proposal@es.net**, BLTierney@es.net

Let us know what we could add/change to make the testbed more useful to you!



# Extra Slides



# Newly Accepted Projects



Project	Summary	Expected Results
<b>Advance Scheduling of Multi-Domain Dynamic Circuits</b> PI: Byrav Ramamurthy, University of Nebraska-Lincoln	This project will investigate multi-domain dynamic circuit creation, and study the issues related to large data transfers over multi-domain circuits.	Better understanding of and multi-domain dynamic circuit creation, and additional features in the existing control plane architecture of OSCARS.
<b>Usability Investigations for High Energy Physics Analysis</b> PI: Ruth Pordes, FNAL	This project will study the issues related to end-to-end integration and use of 100Gigabit networks for the event simulation and analysis applications of physics experiments.	Recommendations to the system administrators, application and middleware developers on changes that would make production use of 100G networks more effective.
<b>Securing Network Services using DASH</b>  PI: Ben Smith, Angel Secure Networks	DASH uses a network of software agents to defend critical software systems from insider and outsider attack and tampering. We will apply this technology to high-performance networks. This project will use the ANI testbed to acquire a better understanding of the type of networks used for DOE science and to demonstrate how our system can help to protect them.	We will demonstrate a network of our software agents (ANGELs) protecting routers on the ANI testbed under a variety of attack scenarios.

# Newly Accepted Projects



Project	Summary	Expected Results
<b>PERT TCP</b> PI: Narasimha Reddy, Texas A&M University	This project will test PERT TCP on a real 10Gbps network over long distances.	A comparison of PERT to TCP-SACK and other high-speed protocols over single flow performance and live video delivery performance.
<b>Scalable Optical Networking with OpenFlow</b> PI: Ben Yoo, UC Davis	This project will design and conduct testbed experiments of OpenFlow based future ESnet. Intelligent and agile network infrastructures with 1) scalable optical networking, 2) hybrid packet/circuit-switched networking, and 3) multi-layer multi-domain network measurement and monitoring will be explored.	This project will investigate, develop, test, and help standardization of OpenFlow towards supporting scalable and dynamic optical networking.
<b>Measuring Energy Efficiency In Networks</b> PI: Thierry E. Klein, Bell Labs / Alcatel-Lucent	The main purpose of this project is to gain an understanding of the power-rate profile and the energy efficiency in real-world routers, switches and networking equipment and to explore opportunities for improving energy efficiency through dynamic management and control (including rate adaptation and sleep modes).	This project will determine the power-rate profiles of the available equipment in the testbed, and conduct experiments turning the equipment on and off to understand the behavior, the state transition times and power consumption during such operations. We will also investigate the implementation of sleep modes and rate adaptation and quantify the obtained energy benefits. (depending on the hardware capabilities and our ability to modify their operations).

4/21/2011

# Current Projects Actively Using the Testbed



Project	Summary	Expected Results
<b>Archstone</b> PI: Tom Lehman, ISI	To dynamically create “slices” of resources across multiple network layers in a vertically integrated manner, so as to generate virtual network topologies. This requires a highly-advanced path computation element which extends the concept of simple path computation to multi-layer, multi-dimensional topologies.	This project will utilize the ANI Testbed to determine design requirements, test alternatives, and evaluate performance of the developed technologies.
<b>FlowBench</b> PI: Prasad Calyam, Ohio Supercomputer Center	To set up different physical topologies in testbed using resources such as NEC Openflow switches, App Hosts, and Monitoring hosts. On these topologies, we will experiment with Openflow and benchmark performance of GridFTP file transfers with enhanced TCP/UDP variants.	The Testbed will be used confirm that our developed technologies will operate as desired with production network equipment, topologies, and configurations.
<b>HNTES</b> PI: Malathi Veeraraghavan, University of Virginia	Hybrid Network Traffic Engineering Software: The purpose of HNTES is to leverage both an IP datagram network and a high-speed optical dynamic circuit network to best serve users' data communication needs.	Experiments on the testbed will be conducted to determine whether flows can be redirected on-the-fly to newly established optical circuits without impacting TCP behavior, and user-perceived performance.

# Current/Future Projects Actively Using the Testbed



Project	Summary	Expected Results
<b>Climate 100</b> PI: Alex Sim, LBL	The Climate100 project integrates massive climate datasets, emerging 100 Gbps networks, and state-of-the-art data transport and management technologies. The goal of this project is to improve the understanding and use of network technologies and transition the climate community to a 100 Gbps network for production and research.	The testbed will be used to test the direct memory access over the network and new data transfers/management algorithms including the use of the 100G transfer protocol.
<b>Projects waiting for 100G prototype Network</b>		
<b>100G FTP</b> PI: Dantong Yu, BNL	This project will design and develop an ultra high speed end-to-end file transfer protocol and tool to move science data at a speed of 100 gigabit per second (Gbps) across the national scale 100Gbps data network interconnecting research centers.	The testbed will be used to verify that this tool scales to 100Gbps on a single wavelength or multiple modulated wavelengths.
<b>100G NIC</b> PI: Jesse Wen, Acadia	This project will develop network interface controller (NIC) hardware and device-driver/protocol-specific software for host and gateway systems operating at 40 and 100 Gb/s.	The testbed will be used to investigate issues that do not arise in initial back-to-back testing. Such issues include interoperability with core-network gear and the effect of long-haul physical impairments.